An Introduction to Non-Cooperative Game Theory
Fall 2005

University of Utah Math Circle
Aaron McDonald

Introduction

We interact with each other and our environments every day. Some of these interactions are predictable, others are not. For those which are not predictable, a number of scenarios with corresponding consequences can occur. For instance, say two siblings Mary and Justin are granted the opportunity to watch television or read their favorite magazines following a hard day of yard work. Assuming that Mary and Justin might choose either of these two opportunities, we have four scenarios to consider. Quite often the interactions we are most curious about are those which involve conflict. Simply put, conflict occurs when two or more of the interacting parties disagree on how an interaction should play out. If you continue to think about Mary and Justin, you should notice that conflict is a distinct possibility if each sibling really wants to watch a different television program on a television that lacks split screen technology. How should Mary and Justin deal with uncertainty in the other sibling's response and conflict? Should Justin decide to watch television knowing that if Mary also chooses to watch television the outcome would be unfavorable for him since his free time would be spent flipping the television between their two programs? Or should he simply pick up the magazine and avoid conflict all together? Mary has the same dilemma. Each sibling wants to make an intelligent decision in the context of the problem, but what is it? Intelligent decision-making is the topic we will discuss over the course of the next two sessions. In particular, we will be interested in understanding what it means to make intelligent decisions and what these decisions are for various games through the use of mathematical modeling.

Outline for October 5, 2005

Section 1.1 Building an Appropriate Model
Section 1.2 Adding Uncertainty

Outline for October 12, 2005

Section 2.1 Intelligent Decision Making
Section 2.2 Best Replies
Section 2.2 Nash Equilibrium

Section 1.1: Building an Appropriate Model

Game theory provides a classic, popular, and quite successful example of mathematical modeling. Generally speaking, mathematical modeling involves the quantification of observations taken from real world phenomenon. These observations can come from thought experiments or from actual experiments conducted. Our goal today is to build a mathematical model that will help us understand the process of making intelligent decisions.

I will begin by introducing fundamental game theory language. Although the words players, actions, strategies, and payoffs might seem straightforward to you, there are detailed definitions for each of these words in game theory. The next few pages are specifically devoted to formalizing these definitions through the use of examples; however, I will briefly cover them now. **Players** are interacting parties that have the ability to display various actions during an interaction. An **action** is a particular physical response which evokes consequences, called

payoffs, for each player involved. A **strategy** is a mathematical description of how players choose an action from those available and an **action profile** is a characterization of the actions that each player has chosen. Lastly, and more formally speaking, a **game** is the complete characterization of an interaction. This characterization must include an explicit description of the set of players, the actions available to each player, and the consequences each player realizes for each action profile the interaction allows.

With this new game theory lingo fresh in our minds, we can reformulate our desires more precisely. Our ultimate goal is to reveal strategies that each player should adopt in a game given some "intelligence" or "rationality" criterion. I have not yet formally stated what it means to make "intelligent" or "rational" decisions. To do so requires that we first develop notation for any game of interest. Upon completion of this task, we will come up with the criterion together. This criterion will enable us to use our model to find solutions (i.e. intelligent decisions) to the game. To begin this process, we start with the very basics: writing down the set of players involved in an interaction and the actions available to each player. Let's begin with a famous example to illustrate this initial step in the characterization of a game.

The Prisoner's Dilemma

Two suspects are being held separately in prison. They are accused of having conspired together to commit a crime that they have previously agreed together to deny if summoned by the police for questioning. A detective is trying relentlessly to persuade each of them to implicate the other with the following inducement. If neither of them confesses they will both be set free because there is not enough evidence against them; if both confess, they will both be punished; but if only one confesses, the confessor is set free and in addition receives a reward while his partner is punished more severely than if he had confessed.

Take a minute to think about the situation. Is conflict a fundamental property of this interaction?

Example 1.1.1 If we assume that each suspect is suspicious of his accomplice's loyalty, describe the emergence of conflict in the Prisoner's Dilemma.

If I honor the pact, then I could get badly burned if my accomplice does not. If I ignore the pact and rat out my accomplice, then I get no jail time plus a nice reward. It is probably reasonable to assume that my accomplice is thinking the same thing about me. What does that mean for me? I know that if we both choose to ignore the pact, I am going to jail. I certainly don't want that to happen. I think that I would like my accomplice to remain quiet and I will ignore the pact. I bet that my accomplice came to the same conclusion about me. As a result, I understand conflict is a fundamental property of this interaction since we each disagree on how the other should behave.

Example 1.1.2 Write the set of players (P) and the set of actions (A_{Player}) available to each player.

Players	$P = \{Suspect\ 1,\ Suspect\ 2\}$
Actions	$A_{Suspect\ 1} = \{Honor,\ Ignore\}$
	$A_{Suspect\ 2} = \{Honor,\ Ignore\}$

The Prisoner's Dilemma is an action−symmetric, 2-player game. We call it this because the interaction involves two distinct participants and each participant has the same action set available to her. For the most part, nomenclature in game theory is straightforward. Let's press on!

An **action profile** is a description of the actions each player has chosen. Since each suspect in the above example has two actions available, there are four action profiles which define how the interaction could play out. For instance, with the Prisoner's Dilemma in mind, one action profile is that each suspect chooses the action Honor in which each honors the pact they made prior to interrogation. The notation we use to write an action profile was chosen to express information in a minimalistic way and is not immediately obvious. A formal definition is given below in the case that our interaction involves n distinct players.

Definition 1.1.1 Action Profile

Assume there are n players participating in a game. An action profile is written as

$$(a_1, a_2, \ldots, a_n)$$

where each component a_i details the action of player i for $i = 1, 2, \ldots, n$.

Notice in the above definition that we have labeled our players numerically. Before we can use the action profile, it is important that we assign each participant to a player number. Sometimes the assignment is clear, other times it is not. For example, in the Prisoner's Dilemma, it is clear that "player 1" is suspect 1 and "player 2" is suspect 2. However, in the case that we label players by their names, the assignment is not clear. No need to worry though; the assignment is arbitrary, but it is an essential detail that you must clearly provide your reader with. We will encounter an example of this in the exercises at the end of this section. In the meantime, let's write out the action profiles for the Prisoner's Dilemma game. Go!

Exercise 1.1.1 What are the four action profiles of the Prisoner's Dilemma game?

Action Profile Specification:

Action Profiles:

If you were to pick up a book on game theory or google it on the web this week (and I encourage you to do this!), you might encounter the **action profile space** (or pure strategy space). The action profile space is simply the collection of all action profiles a game allows. For instance, the collection of the four action profiles above forms the action profile space of the Prisoner's Dilemma game. Those of you interested in set operations may enjoy a more formal definition, and so I provide it below.

Definition 1.1.2 Action Profile Space

Suppose we were to label each player's action set by A_i, where the subscript i denotes player i as usual and $i = 1, 2, \ldots, n$, then the action profile space for the game is given by $A = A_1 \times A_2 \times \ldots \times A_n$.

The next step in building a model is to think about and write down the payoffs each player receives given some action profile. What does this entail? In the Prisoner's Dilemma, we remember that the action profile (Honor, Honor) yields each player a particular payoff and the action profile (Honor, Ignore) yields each player a different payoff. This clearly implies that the payoff each player receives *depends* on (or is the consequence of) the action profile. This means that the payoff each player receives is a *function* of action profile.

Definition 1.1.3 Payoff Function for Player i, $\pi_i : A \to \mathbb{R}$

The payoff function for player i maps each action profile in the action profile space to the payoff that player i receives from this action profile.

To provide the payoff function for each player requires that you explicitly write a function for your reader. The general idea of payoff relationships to a player for each action profile should be given in the verbal description of the game. Quite often this is the case. So now what? It becomes your duty to assign the numerical values for the payoffs taking into consideration that these payoffs must correspond with the problem. I will help you get started.

Exercise 1.1.2 Consider the Prisoner's Dilemma (Example 1.1.1). Suppose that suspect 1 is player 1 and suspect 2 is player 2. Use a to denote any action profile ($a \in A$). Construct the payoff function for each player.

a. Construction of the payoff function for player 1.

b. Construction of the payoff function for player 2.

It is quite common in this business to express all of the above ideas (players, actions, and payoffs) for 2-player games visually in the form of a matrix. In particular, the intersection of rows and columns of the matrix make up action profiles and entries assigned at these intersections correspond to the payoffs to each participant. This matrix is called the **payoff matrix**. Formal descriptions are usually litered with indices. We will bypass a formal definition and simply learn by doing. Below is the general plan for the construction of a payoff matrix for a 2-player game in which each player has two actions available to her. For the purpose of demonstration, label the two actions available to player 1 $\dot{a}_1$ and $\bar{a}_1$ and the two actions available to player 2 $\dot{a}_2$ and $\bar{a}_2$.

$$\begin{array}{cc} & \text{player 2} \\ & \begin{array}{cc} \text{action } \dot{a}_2 & \text{action } \bar{a}_2 \end{array} \\ \text{player 1} \begin{array}{c} \text{action } \dot{a}_1 \\ \text{action } \bar{a}_1 \end{array} & \left(\begin{array}{c|c} (\pi_1(\dot{a}_1,\dot{a}_2),\ \pi_2(\dot{a}_1,\dot{a}_2)) & (\pi_1(\dot{a}_1,\bar{a}_2),\ \pi_2(\dot{a}_1,\bar{a}_2)) \\ \hline (\pi_1(\bar{a}_1,\dot{a}_2),\ \pi_2(\bar{a}_1,\dot{a}_2)) & (\pi_1(\bar{a}_1,\bar{a}_2),\ \pi_2(\bar{a}_1,\bar{a}_2)) \end{array}\right) \end{array}$$

Exercise 1.1.3 Construct the payoff matrix for the Prisoner's Dilemma game using the payoff functions you defined in Exercise 1.1.2.

Section 1.1 Problems:

Let's take some time to think about a new problem. I will give you another classic game. This one is typically entitled the Battle of the Sexes. However, the basic storyline does not rely on the fact that the 2 players are of different sex. I have made up my version of the story and titled it the Battle of the Friends.

Directions: Read the following verbal description of the interaction and answer the questions below.

Battle of the Friends

Adam and Bob want to catch a concert together this weekend. After a little research online, they discovered that two concerts still have seats available: one by 50-Cent and one by George Winston. Bob prefers to see 50-Cent while Adam prefers to see George Winston. If they go to different concerts, each of them is equally unhappy listening to the music of the other artist alone.

Problem 1 What is the set of players for this game?

Problem 2 What actions are available to each player?

Problem 3 What are the four possible action profiles? Your solution must include an action profile specification.

Problem 4 Construct a payoff function for each player.

Problem 5 Construct a payoff matrix that characterizes the above concepts.

Section 1.2: Adding Uncertainty

So far we have covered the basic construction of a model to analyze conflict. This included the complete description of the set of players, the actions available to each player, and the payoffs associated with each action profile. However, we have not dealt with the complex manner in which players often play games. Specifically, over the course of many games, players are often unpredicatable in their action choices. That is to say that players are not consistent in their action choice from match-up to match-up. Players entering a game are aware of this inconsistency, or uncertainty, in their opponent's response. Recall that one of our goals was to build a model that accounted for this uncertainty. Our current model lacks this structure. How do we incorporate uncertainty into our model? In mathematics, probability is the language of uncertainty. In order to model uncertainty in opponent action choice, a short introduction to probability theory is required.

Assume we are interested in some **experiment** that has more than one possible **outcome**. Some outcomes may be more likely than others. Consequently, we might like to have a measure of how likely a particular outcome is. This will provide us with a notion of how often this outcome would occur if we were to perform the experiment numerous times. This is the basic idea of probability theory. A **probability** is a measure of the frequency in which a particular outcome for our experiment would occur if we repeated the experiment an infinite number of times. Probabilities are usually expressed as fractions that detail the fraction of time we expect to observe a particular outcome. Mathematically, this means that a probability is a real number on the interval [0,1]. We also can assert that the sum of the probabilities of all the outcomes of an experiment must be 1. These two conditions are fundamental probability axioms.

But how can we get our hands on probabilities in general? Often we can hypothesize the probabilities of outcomes of an experiment so that we don't actually have to run the experiment forever. The two examples below will help demonstrate these ideas.

Example 1.2.1 Flipping a Fair Coin

Experiment: Flip a fair coin.

Outcomes: There are two possible outcomes that may result from flipping a fair coin: heads and tails.

Probabilities: $Pr(Heads) = \frac{1}{2}$ and $Pr(Tails) = \frac{1}{2}$

Notice that $0 \leq Pr(Heads) \leq 1$, $0 \leq Pr(Tails) \leq 1$, and $Pr(Heads) + Pr(Tails) = 1$.

Exercise 1.2.1 Rolling a Six-Sided Die

Experiment: Roll a Six-Sided Die

Outcomes:

Probabilities:

You might be wondering how this applies to our problem. Remember that we need to model uncertainty in player action choice. Using the above examples as inspiration, take a minute to think about the "experiment" we are conducting and the possible "outcomes" for this experiment.

EXPERIMENT:

OUTCOMES:

To accommodate the fact that each player may choose to assign probabilities to these outcomes is different manners, we will simply assign each probability a constant and assume collectively the probabilities adhere to the probability axioms previously discussed. This assignment is called a **probability density function** for our experiment. It is a function because it maps each experimental outcome (the actions) to the probability of that outcome.

Example 1.2.2 Construct the general probability density function for player 1 of the Prisoner's Dilemma.

Let x represent any outcome for the experiment of interest and call $pdf_1(x)$ the probability density function for player 1 of the Prisoner's Dilemma. We can explicitly write the function as

$$pdf_1(x) = \begin{cases} x_{11} & if\ x\ =\ Honor \\ x_{12} & if\ x\ =\ Ignore \\ 0 & otherwise \end{cases}$$

subject to the conditions $0 \leq x_{11} \leq 1$, $0 \leq x_{12} \leq 1$, and $x_{11} + x_{22} = 1$.

Exercise 1.2.2 What does the assignment $x_{11} = \frac{1}{3}$ and $x_{12} = \frac{2}{3}$ mean?

Exercise 1.2.3 Construct the general probability density function for player 2 of the Prisoner's Dilemma. Do so in the spirit of Example 1.2.2.

We are finally ready for a formal introduction to strategies. You may recall earlier that I defined a strategy as a mathematical description of how players choose an action from those available. That mathematical description is in fact a probability density function over an action set.

Definition 1.2.1 Strategy

A strategy is any valid probability density function over a well-defined action set.

There are two types of strategies which require discussion: pure and mixed strategies. A pure strategy is a strategy in which all probability is assigned to one action. This strategy is predictable since a player who adopts this strategy always utilizes the same action. All other strategies are called mixed strategies. These strategies necessarily have non-zero probability assigned to two or more actions. You should think of these strategies strategies as a mixing of two or more actions which ultimately results in unpredictability in action choice.

Exercise 1.2.4 What are the two pure strategies that player 1 of the Prisoner's Dilemma has available to her?

Terrific. Now let's make sure we got mixed strategies down.

Exercise 1.2.5 Write two distinct mixed strategies that player 1 of the Prisoner's Dilemma has available to her. Note that there are a bunch of them, but just choose any two you like best.

Exercise 1.2.6 Write down a strategy (pure or mixed) for each player of the Prisoner's Dilemma.

So how do we compute payoffs when players utilize strategies? We know that in the case that each player uses a pure strategy, only one action profile can result and these payoffs are already defined. However, if at least one player uses a mixed strategy, the computation becomes difficult since multiple action profiles are now a distinct possibility. What now? The best we can hope for is to determine the average payoff each player receives from such a situation. We begin by thinking about repeating the game many, many times. In doing so, we would like to determine how often particular action profiles are used by the players collectively. Mathematically, this means we would like to calculate the probability of each action profile given the strategies of each player. This information is useful because it will allow us to compute a weighted average of the payoffs (called the **expected payoff**). This calculation then provides us with an expected payoff that a player shall receive from engaging in a single interaction.

Definition 1.2.2 Expected Payoff for Player i

The expected payoff that player i shall receive from an interaction is given by $u_i = \sum_{a \in A} Pr(a) * \pi_i(a)$ where Pr(a) is determined from the strategies $pdf_1, pdf_2, \ldots, pdf_n$.

Wow. That seems kind of complicated. Let's do an example so that you can see how it works. Consider the Prisoner's Dilemma game with any appropriate payoff matrix.

Assume that the two players are using the strategies given below:

$$pdf_1(x) = \begin{cases} \frac{1}{3} & if\ x\ =\ Honor \\ \frac{2}{3} & if\ x\ =\ Ignore \\ 0 & otherwise \end{cases}$$

and

$$pdf_2(x) = \begin{cases} \frac{3}{5} & if\ x\ =\ Honor \\ \frac{2}{5} & if\ x\ =\ Ignore \\ 0 & otherwise \end{cases}$$

Recall that the action profile space for this game is given by

$$A = \{(Honor, Honor), (Honor, Ignore), (Ignore, Honor), (Ignore, Ignore)\}$$

By definition, the expected payoff player 1 shall receive is given by

$$\sum_{a \in A} Pr(a) * \pi_1(a) \;=\; Pr((Honor, Honor) * \pi_1((Honor, Honor)) + \ldots + Pr((Ignore, Ignore)) * \pi_1((Ignore, Ignore)).$$

Now is a great time to discuss a major issue that we have bypassed. You should notice that in the above statement, we must compute the probability of each action profile. This is problematic! Recall that the strategies pdf_1 and pdf_2 tell us how often we can expect each player *individually* to chose a particular action. However, these strategies do not tell us exactly how players choose actions *collectively* to form particular action profiles. That is precisely what we need to resolve in order to continue. Let's use our intuition to figure it out.

Example 1.2.3 Calculate $Pr((Honor, Honor))$ given the mixed strategies pdf_1 and pdf_2.

We know that player 1 will choose Honor $\frac{1}{3}$ of the time and player 2 will choose Honor $\frac{3}{5}$ of the time. Assuming each player chooses their action *independent* of the other players according to their strategy respectively, we might guess that the action profile (Honor, Honor) should occur $\frac{1}{3} * \frac{3}{5}$ of the time. This is exactly right! Then, $Pr((Honor, Honor)) = \frac{3}{15}$.

Nice work. Let's use this idea to calculate the probability of the other action profiles.

Exercise 1.2.7 Calculate $Pr((Honor, Ignore))$, $Pr((Ignore, Honor))$, and $Pr((Ignore, Ignore))$ given pdf_1 and pdf_2.

Exercise 1.2.8 Calculate the expected payoff that player 1 of the Prisoner's Dilemma shall receive when player 1 uses the mixed strategy pdf_1, player 2 uses the mixed strategy pdf_2, and all payoffs are given by the payoff matrix below.

	Honor	Ignore
Honor	(0, 0)	(−3, 3)
Ignore	(3, −3)	(−1, −1)

Exercise 1.2.9 What does the above result mean?

Section 1.2 Problems:

Problem 1 Calculate the expected payoff that player 2 shall receive from the Prisoner's Dilemma using the same information given in Exercise 1.2.8.

Problem 2 Using your results from Exercise 1.2.8 and Problem 1 above, can you predict who is happier (on average) with their payoff?

Problem 3 Fix pdf_1 and find another strategy for player 2 that yields a greater expected payoff to player 2 than pdf_2 does.

Problem 4 Can you find a strategy for player 2 which yields an expected payoff to player 2 that is at least as good as any other strategy can given that player 1 uses the strategy pdf_1? Label this strategy for player 2 $\bar{pdf}_2$.

Problem 5 Can you find a strategy for player 1 which yields an expected payoff to player 1 that is at least as good as any other strategy can given that player 2 uses the strategy $\tilde{pdf}_2$?